Who Can Tell If It Is AI: Can Inductive Learning Improve AI-Art Literacy?

Dávid Csomor

S4304837

Department of Psychology, University of Groningen

PSB3E-BT15: Bachelor Thesis

Group number: 2324_2a_45

Supervisor: Ben Gützkow

Second evaluator:

In collaboration with: Jona Frank; Timo Landzettel; Paul Rühsen; Leonard Schmidt; and Axel Warbroek

June 20, 2024

Abstract

AI-generated art has been undergoing a rapid development in the past years both in terms of quality and efficiency. The aim of the current study was to investigate how good people are at distinguishing AI from non-AI art; and whether this ability can be improved by training. The method of choice for the training was based on the inductive learning paradigm, which prioritizes intuitive pattern recognition, rather than learning based on factual information. Thus, participants in the experimental group were instructed to observe artworks labelled as AI or non-AI; and, subsequently, they were asked to complete a test where they could label the artworks as AI or non-AI themselves. The results showed that participants performed at chance level in the absence of training; while participants in the experimental group performed only slightly, but statistically significantly, better. Additional variables were also explored as potential influences on test performance, such as confidence; AI/non-AI art knowledge; and the types of the artworks. The results suggest that, while AI art may be indistinguishable from non-AI art at first glance, people may be able to develop a better intuition for AI art recognition. The study was limited in statistical power though, and the homogeneity of the sample also warrants more studies to establish better generalizability.

Who Can Tell If It Is AI: Can Inductive Learning Improve AI-Art Literacy?

In the past years, the importance and influence of Artificial Intelligence (AI) has been growing rapidly, due to its recent, substantial achievements. For instance, AI is now capable of producing coherent text, even pieces of literature and scientific articles. It can also generate computer codes with high accuracy, and much faster than humans. AI can already generate images of people that do not exist but seem quite real, and it can also generate images of people that do exist but are placed within a different context. The phenomenon of *AI hyperrealism* demonstrates just how real these images can seem: a recent study showed that AI generated faces were judged to be real more often than the faces of real humans (Miller et al., 2023). If these difficulties would arise in everyday life more routinely, it could facilitate deception and propaganda on multiple forms of media, including social media. For instance, AI images could be used as fake evidence depicting politicians performing a scandalous action; or being at an inappropriate place. The risk could be especially high in places where media is much more controlled, and thus where there is no way to counterbalance fake news, or where AI-literacy is lower. Thus, developing sensitivity for AI generated content may be highly relevant to the coming years.

There are also impressive technological advances in the domain of AI art. In comparison to other domains, art is widely seen as a uniquely human capacity, as a recent study has suggested as well (Bellaiche et al., 2023). However, developments in the past few years may also call this into question. In 2023, Boris Eldagsen submitted an AI generated piece to the Sony World Photography Awards; where the contestants were meant to be human photographers. Boris succeeded in winning the award; which he decided to reject, as the aim of his experiment was not to take the prize, but to spark a rather scandalous debate. The scope of AI-art[1] extends to paintings as well: advanced tools, such as Midjourney, can already generate digital paintings that look quite realistic: a previous study has suggested that the accuracy of people's guesses about whether a piece is an AI or non-AI artwork is not better than random chance level, even though they were more likely to prefer the aesthetics of non-AI art (Zhou et al., 2023).

[1] Throughout the text, I will refer to art generated using AI tools, such as Dall-E or Midjourney, as AI-art; and to art fully made by human hands, without such tools, as non-AI art.

These phenomena can cause serious copyright issues as well though, since people can generate almost literal copies of developed styles of artists in a matter of seconds, potentially generating not only artworks, but considerable financial gains as well: the AI art of the highest-value ever sold on an auction was worth as much as $432,000 to its buyer (Academy of Animated Art, n.d.). As the algorithm that generated this artwork was trained on 15.000 different portraits, it is not traceable back to any specific one of them, which makes it virtually impossible to raise plagiarism charges.

Looking at these tendencies, it is crucial that people are aware of how dominant AI is becoming in the art industry and beyond, and that they are equipped with skills that can help them distinguish between AI and non-AI content. In addition, it is also important to assess how much difficulty people may already have with making these distinctions, and whether they can learn to notice if a certain content was generated by AI. AI is a new technology; thus part of the reason for the current difficulty of recognizing AI content could be that people just have not been exposed to it enough yet. More exposure, in combination with some guidance, may be a potential solution to this issue. To explore the extent of this potential, the current study will test a training method which will attempt to increase the accuracy of people's judgments. In addition, it will also investigate the general capacity of people to distinguish between AI-artworks and non-AI artworks.

A key assumption of this study is that AI-art recognition can be seen as a skill that can be improved. We argue that inductive learning may prove to be a useful method for this. Inductive learning is a largely intuitive process which involves acquiring knowledge about certain phenomena through observing them (Kornell & Bjork, 2008). In the context of object recognition, inductive learning would occur if someone observed exemplars and recognized certain patterns of similarities and differences between them. This recognition does not necessarily have to be conscious though. Even if someone cannot describe the patterns that they notice, their abilities of object recognition and categorization can still improve (Kornell & Bjork, 2008). This phenomenon can be relevant to the recognition of artwork styles as well: there is evidence that humans can intuitively learn to distinguish artworks by different artists, even if their styles are reasonably similar (Kornell & Bjork, 2008; Kang

& Pashler, 2012). Thus, we argue that inductive learning can be suitable for our research purposes, as the clear differences between AI and non-AI content are ill-defined.

There are an explicit set of features which are often used today to tell if images are AI-generated (e.g., too many fingers on a hand, nonsensical text), but these are applicable only to a subset of cases. Furthermore, even these differences that are currently explicitly identifiable might be subject to change as AI algorithms are going to improve. Therefore, since the differences between AI and non-AI art are difficult to define or even to consciously perceive, developing a better intuition indeed seems to be the most plausible way to differentiate between them with higher accuracy. Even if the distinguishing features would be so subtle that they could not be verbalized, it would still be possible to develop a better gut feeling for recognizing AI-art. In general, the importance of intuitive decision making has already gained support by previous research both in terms of its effectiveness and its usefulness on a day-to-day basis (Sedimeier, 2014; Betsch and Glöckner, 2010). In addition, in case there are more definable, unique patterns that are typical to AI art, inductive learning can also lead to becoming consciously aware of those patterns, even if, initially, they are perceived only unconsciously – something that can also occur when people learn a new language through inductive learning methods (Abdukarimova & Zubaydova, 2021). People can be exposed to words and sentences of a language they are currently learning, they can notice certain grammatical structures within them, which they can also become able to explicitly formulate themselves. Potentially, the same may occur when people observe AI and non-AI art as well: after lots of structured exposure to AI art, people may gain insights to how can they identify AI art. In turn, this may also pave the way for a more general understanding of such patterns – although the current study does not extend its research questions towards these directions.

Inductive learning can also be divided into more specific subcategories based on how the exemplars are presented to the learner. The stimuli can be 'massed' – meaning that the instances of the same category can be presented to the learner without mixing the different types of stimuli while presenting them. In contrast, another way to present exemplars is called „interleaving" – in which stimuli from different categories are mixed together while being presented. This can be done by presenting more stimuli either all at once; or only one at a time – in the latter case, the stimuli from one category would be followed by a stimuli from a different category. Interleaving has been shown by previous research to be more advantageous in scenarios where the stimuli are more complex; and where the categorical differences are subtler, and thus more difficult to notice (Kang and Pashler, 2012). This advantage may be due to perceptual discrimination, which relies on perceiving the differences between categories as contrasts between them. If stimuli of different categories are perceived in contrast with each other, the subtle differences between them can become more perceptible and/or identifiable more easily (Birnbaum et al., 2013). Presumably, interleaving has more ecological validity as well: in real life, people are probably exposed to stimuli from multiple different categories rather than to stimuli from the same category at a time. In the context of AI art, if someone is scrolling on a social media platform, it is unlikely that the feed will present only AI or only non-AI images in succession. Thus, interleaved inductive learning may be the best way to assess whether people can indeed learn to appreciate the qualitative differences between AI and non-AI art.

In the final part of the introduction, we present the hypotheses for the current study. First, in line with the previous explanations, our first hypothesis predicts significantly more accurate test performance for participants who complete the training procedure. We have no specific predictions with regards to how participants will perform in the control group; but we will explore those results as well in our analyses.

Hypothesis 1: People will become better at judging whether an artwork is AI or non-AI after structured exposure, based on interleaved inductive learning.

Next, we aim to explore additional idiosyncrasies that may be related to the recognition of AI art, regardless of the training the participants receive. We argue that knowledge in either the AI or non-AI art domain can significantly improve performance. Simply put, those who know more about art may have a better understanding of the subtle details that characterize human-made art. Likewise, people who have more knowledge about or experience with art-generating AI tools, might also be better at judging the images, since they likely have more exposure to AI art; which may have already caused a form of inductive learning prior to the study.

Hypothesis 2: More knowledge about either AI art or non-AI art will predict significantly better test performance in the whole sample.

In addition, inspired by the work of Miller and colleagues (2023), we will also investigate how confidence about the judgments might be related to the actual outcome. Miller has found that higher confidence predicts worse performance in the recognition of AI generated faces; although the study does not include an explanation of why this might be the case (Miller et al., 2023).

Hypothesis 3: Confidence about test performance across the stimuli will not significantly predict actual test performance.

Finally, we also explore whether the type of the presented artwork has an influence on the accuracy of judgement. Kornell and Bjork (2008) exclusively used landscapes in their inductive learning study. We explore whether abstract art, and portrait art – in addition to landscapes – show a different pattern, i.e., whether they vary in difficulty. In addition, we will also test whether there will be a difference in accuracy of judgement between AI and non-AI artworks. The analyses will take into account these categories, but will be exploratory; no specific hypotheses are defined about them. We include these categories to stimulate future research and theorizing in this area.

Methods

Participants

On one part, the sample (*N* = 100) contained participants collected via the SONA-systems platform from first-year psychology students at the University of Groningen (*n* = 35), who received course credits for their participation. The other participants were recruited through convenience sampling based on the social network of the authors. Data cleaning excluded 18 participants who gave insufficient answers (i.e. below 20). The final sample used in this study therefore consisted of 82 participants. No demographic data was recorded.

Design of the Stimuli

A set of 120 images was compiled, consisting of 60 AI-generated pictures and 60 traditional artworks. The AI-generated artworks were created with the software package Midjourney (2023) during March 2024. An example of a prompt is [/imagine old renaissance portrait of a 14th century peasant] or [/imagine oil on canvas landscape after sundown, with a vibrant, purple, but still realistic sky, depicting a slightly hilly, clean but picturesque field. in the style of Herman van Swanevelt]. A full list of prompts is in Appendix A. Through this process pictures were created in three categories: abstract art, portraits, and landscape art. Twenty pictures were selected for each category, equalling a total of 60 AI-generated images. This selection was made by voting among the researchers, on the basis that the selected pictures should fulfil the following requirements: they should not be easily identifiable as AI-generated images, and there should be some variety within the respective categories.

The non-AI artworks were selected from a variety of sources. Most of the images were sourced from the website of the Metropolitan Museum of Art (http://metmuseum.org), while some additional images were found from other websites. Again, we opted for 20 pictures from each of the previously mentioned categories.

Procedure & Measures

The participants were asked to complete the study online, on the platform Qualtrics. At the start of the experiment, the participants were asked to fill out the questionnaires about Art knowledge and about AI Art interest and affiliation, which were adapted from the Vienna Art Interest and Art Questionnaire Knowledge (VAIAK; Specker et al., 2020).

Art knowledge

For the assessment of art interest, we used a 7-item scale based on Specker and colleagues' (2020) Vienna Art Interest and Art Knowledge Questionnaire (VAIAK). Artistic interest was measured across two scales, with four items capturing self-reported interest rated on a 7-point Likert scale (1 = *not at all*, 7 = *very much*) and three behavioural items rated on a 7-point frequency scale (1 = *less than once per year*; 7 = *once per week or more often*). The self-reported art interest scale included items such as: "I am interested in art" and "I am always looking for new artistic impressions and experiences". Examples of the behavioural items are: "How often do you visit art museums and/or galleries?" and "How often do you read books, magazines or catalogues about art?". The internal consistency of the artistic interest scale that was used in this study was good (α = 0.86).

AI Art affiliation

For the assessment of AI interest, we adapted the VAIAK scale (Specker et al., 2020) to ask about AI image generation instead. We adapted the items in such a way that the new scale measures self-reported AI interest using four items rated on a 7-point Likert scale (1 = *not at all*, 7 = *very much*) and three behavioural items regarding AI rated on a 7-point frequency scale (1 = *less than once per year*; 7 = *once per week or more often*). The self-reported AI interest scale included items such as: "I am interested in AI art technology" and "I like to talk about AI art technology with others". Examples of the behavioural items are: "I´m always looking for new AI art Impressions and experiences?" and "How often do you seek out AI art technology?". The internal consistency of the AI interest scale that was used in this study was good (α = 0.801).

After the completion of these questionnaires, participants were given the instructions for the experiment itself. The experimental group and the control group were given partially different instructions, as the experimental group was asked to complete both a training and a testing procedure, while the control group was only asked to complete the testing procedure. However, the testing procedure was identical for both groups.

The experimental group was first asked to observe the artworks that appeared on the screen. Then, the artworks were shown, each with a label indicating whether the artwork is AI or non-AI. Each artwork was shown for a duration of 5 seconds; with 2 seconds of break in between the stimuli. In total 78 artworks were shown, of which 39 were AI and 39 were non-AI. Within the AI and non-AI-pool each, 13 portrait artworks, 13 landscape artworks, and 13 abstract artworks were presented. The order of the presentation followed the interleaved spaced design of inductive learning (Kang & Pasher, 2011). An AI artwork was always followed by a non-AI artwork, and *vice versa.* After all the artworks were shown, the training part of the experiment was over. Participants in the experimental condition were able to take a short break and continue with the testing phase.

In the testing part of the experiment, all participants were asked to guess whether the artworks they were presented with one by one, another set of 42 artworks, were AI or non-AI.

Image Classification

The classification of images as AI-art or human-art was captured with a single item: “Was this artwork made by a human or by Artificial Intelligence (AI)?”. There were two response options (“Human-made” or “AI-made”). Participant’s confidence in their classification was also assessed using a single item asking: “How certain are you in your judgment?” on a slider from 0 to 100.

They were also asked to indicate how much they liked each artwork; a Likert-scale was applied. Each artwork was presented together with the two scales. Like in the training set, the pool contained an equal number of artworks from each subcategory; but it consisted of a different set of artworks. After participants in the experimental group were finished with the test, they were asked to write any remark or feedback about the experiment if they wished to. Finally, they could see a message thanking their participation, which marked the end of the procedure.

Results

First, the main effect of training on performance was investigated, as described in Hypothesis 1. In Table 1 (Appendix A), performances for both the control and the experimental group are presented. The mean percentage scores indicate that participants in the experimental group tended to perform better: their guesses whether an artwork is AI or non-AI were, on average, about 6.5% more likely to be correct (M = 57.35%, SD = 10.25) compared to participants in the control condition (M = 50.85%; SD = 9.4). Interestingly, a one-sample t-test showed that participants in the control condition did not perform significantly better than random chance level, $t(46) = 0.623$, $p = 0.54$. A one-way ANOVA analysis showed that the difference between conditions was statistically significant, $F(1, 80) = 8.853$, $p = 0.004$. This means that training improved performance significantly above chance level. Thus, Hypothesis 1 is supported by the analysis.

Next, Hypothesis 2 was tested, which predicted that both AI-art and non-AI art knowledge would be significantly positive predictors of test performance. Multiple regression analysis showed that neither AI-art knowledge ($t(78) = -0.272$, $p = 0.787$) nor art knowledge ($t(78) = -0.787$, $p = 0.434$) was significantly related to test performance (see the regression coefficients in Table 2 in Appendix A).Thus, these predictions were not supported by the analysis.

Hypothesis 3 was that confidence would not be significantly related to test performance. This was indeed the case, as simple linear regression analysis showed (t = --0.324, $p = 0.747$), (see Table 3 in Appendix A for the summary of the results).

Finally, we conducted our exploratory analyses—we looked at how the different types of images influenced test performance. For the tests comparing the non-AI artworks with the AI-artworks in the two groups and in the whole sample, we applied Bonferroni Correction to adjust the alpha level, which became 0.01667, after adjusting for 3 comparisons.

Paired Samples T-Test has shown that performance was not significantly different when comparing non-AI artworks with AI artworks ($t(81) = -0.410$, $p = 0.683$). However, when we looked only at the performance of the experimental group, a paired-samples t-test showed that when the artwork was AI, the guesses tended to be significantly more accurate compared to when they were guessing non-AI artworks ($t(34) = 2.652$, $p = 0.012$). In contrast, the performance of the control group did not depend significantly on whether the artwork was non-AI or AI (t(46) = 1.499, $p = 0.141$). Thus, improvement in test performance due to training only occurred for AI artworks, but not for non-AI artworks.

We also explored how participants tended to perform with portrait, landscape, and abstract artworks in the whole sample. We applied Bonferroni Correction to these tests as well to adjust the alpha level, which became 0.01667, adjusting for 3 comparisons. The mean and standard deviation scores are visible in Table 4 (Appendix A). On average, guesses were the most accurate for portraits and the least accurate for abstract artworks. Participants were significantly less likely to make a correct guess when the painting was abstract, compared to when they were making a guess for portraits ($t(81) = 3.343$, $p = 0.001$). When we compared accuracy for abstract artworks to landscapes, it was not significant under the adjusted alpha level ($t(81) = 2.381$, $p = 0.020$). There was no significant difference between the guessing accuracies for portraits and landscapes ($t(81) = 1.639$, $p = 0.105$). These results, in general, point to a difference in the ability to recognize certain types of artworks as AI or non-AI with better accuracy than others; with abstract artworks being the least distinguishable.

Discussion

The current study aimed to collect evidence about how good people are at guessing whether artworks were made with AI tools or only by the hands of a human, and whether this can be improved by training in the form of inductive learning, with an interleaved arrangement of the stimuli. In addition, the study also explored how the accuracy of people's perceptions of the artworks' origins may be influenced by the type of artwork that they were guessing. About these effects, we did not hypothesize anything specific, but we did suspect nevertheless that the different types may have different features that AI may have a harder time replicating; making it easier for people to notice when an artwork was made by AI or not. We also looked at other potential correlates of guessing accuracy, such as AI/non-AI art knowledge and confidence about judgement; and we investigated whether a discrepancy may be subconsciously detected and expressed in the form of preference for either AI or non-AI artworks.

Discussion of Results for Hypothesis 1

Training did result in significant improvement in judgement – thereby supporting the idea that people's intuitions about the origins of artworks can improve just by exposing them to AI and non-AI artworks and telling them which one is which. As far as we know, this phenomenon has not been researched before, which makes this result a new addition to the scientific understanding of AI art perception. It appears that humans are able to perceive patterns that set AI art apart from non-AI art, and that these patterns get integrated into their implicit knowledge-base; which results in improved recognition of the artworks. Moreover, this suggests that there is another phenomenon for which this study provides evidence. The inductive learning paradigm can only be scientifically meaningful if there are some meaningful patterns to recognize; thus, the results, albeit indirectly, may also provide evidence for the existence of such patterns. In addition, the study may also suggest where these patterns are to be found. One of the significant exploratory results indicated that people did become better at recognizing AI artworks to be AI, but there was no improvement in recognizing non-AI artworks to be non-AI after the training. These results may suggest that AI art might possess certain features that encompass ‘AI-ness’; while non-AI artworks, crafted by the hands of a human, may not have a commonly essential ‘human-ness’ to them – or, at least, people may lack the ability of perceiving it. If this is indeed the case, the differences in the test performances for AI and non-AI art are in line with theories about object categorization, including exemplar theory and prototype theory. These theories, although being different in other ways, share the idea that categorization requires perceptible similarities between the objects in the given category (Dopkins & Gleason, 1997). When it comes to the current study, we may see the judgment of artworks as a form of categorization to either the AI or non-AI ‘category’. Since the categorization of AI artworks was more successful after training, while the same was not true for non-AI artworks, it might be the case that AI art has more perceptible similarities based on which it can be categorized, compared to non-AI art.

Discussion of Results for Hypothesis 2

Next, we discuss the results for the additional hypotheses. First, art knowledge/AI art knowledge did not predict test performance, which does not support our hypothesis. It appears that, in our sample, knowledge in either of the domains was insufficient to assist the participants in making accurate judgments about the artworks. However, this evidence lacks generalizability, as the sample did not contain individuals with very high knowledge in either of the two areas. The differences in the scores would probably have to be higher to explain any variance in the performance – the subtle differences between the artworks probably require extensive education and experience in the art domain. Indeed, a previous study has shown that experts tend to be more accurate at distinguishing AI from non-AI art (Ha et al., 2024). Nevertheless, the lack of any significant influence of non-AI/AI art knowledge on test performance in our sample still suggests that the recognition of “AI-ness” is unlikely to occur based on a moderate degree of factual knowledge in these domains. This may further support the notion though that the average person might do better if they rely on intuition based on prior experience.

Discussion of Results for Hypothesis 3

The prediction that confidence in the participants' own judgements will not be related to test performance was supported by the results, as there was no significant correlation between these constructs. This result is in line with a previous study (Miller, 2023). One potential explanation is that higher confidence could have occurred when people perceived a certain pattern, which they could have attributed to 'AI-ness' or 'non-AI-ness'. Although the results could be biased, since not all participants who completed the tests answered all of the questions, the results were not even close to achieving significance. Thus it is unlikely that confidence could be a meaningful predictor.

Discussion of Exploratory Results

Looking at the differences in test performance for the types of artworks, the only detected significant difference was between portraits and abstract art; although the difference between abstract and landscape artworks would have been significant at the alpha level of 0.05, without Bonferroni correction. These results may be explained by previously mentioned theories of categorization and the inductive learning paradigm as well. Abstract art is abstract precisely because it tends to lack the clear structure that other types of art have; since it is composed of mainly non-figurative elements. Therefore, it is possible that abstract artworks are more differentiated from each other; and thereby it may be more difficult to perceive any systematic similarities between them, that could be attributed to AI-ness or non-AI-ness.

Finally, the results also indicated that participants who did not receive training could not guess the origins of the artworks more accurately than chance level. In line with a previous study as well (Zhou et al., 2023), this study further supports the idea that AI art can be indistinguishable from human artworks, at least at first glance. However, it is important to note that we selected the artworks that served as the stimuli, and we intentionally did not pick artworks that were made by the most recognized artists. We tried to match the quality of the non-AI artworks with those of the AI artworks, since the aim of the study was not to investigate whether AI can reproduce the 'best' art. Rather, the study focussed on whether it is possible to train the ability of humans to tell whether any artwork, regardless of its perceived artistic or aesthetic value, is made by AI or not. According to these results, training may indeed be effective; even if people have no prior skills at all to tell AI apart from non-AI.

Limitations

The aforementioned conclusions have to be considered with caution though, as the study had some shortcomings in terms of its quality, power, and generalizability. First, the sample mainly included Psychology students from the University of Groningen; and from our own social networks – people in different age groups, or with different cultural and educational backgrounds might have displayed different tendencies in AI-art recognition. The size of the sample was also too limited to arrive at robust conclusions. Moreover, some aspects of the study design may have distorted the results as well. On one hand, participants only had 10 seconds to indicate their choices and to answer other additional questions; due to this, many of them could not always complete their responses. On the other hand, the stimuli for the study were chosen by ourselves, and we might have been unconsciously biased towards choosing more differentiated artworks in some artwork categories or subcategories than in the others. Moreover, all the stimuli were generated by a single AI tool, Midjourney; thus the AI artworks might have been more similar to each other. Perhaps it is not AI in general, but Midjourney, that has its own recognizable style; and the patterns participants perceived may have been characteristics of this style. Finally, it should be noted that the improvement in test performance was not too substantial; but we argue that it still has practical significance. The duration of the learning process was relatively short - perhaps a study including exposure to more stimuli in the training condition, over a longer time span, may lead to more substantial improvements.

Overall Implications and Future Directions

The current study has a number of implications for the domain of cognitive psychology, while also being relevant for the artistic world. When it comes to the relevance for cognition, the study supports the efficacy of intuitive learning processes, which may become increasingly relevant in the coming years and decades. As AI is already generating a vast amount of coherent text, images, or even videos that seem ‘real’, we may encounter a vast number digital artefacts that can potentially be misleading, or even intentionally deceptive. In turn, the recognition of subtle patterns may become of prime importance to human intelligence, perhaps more than ever. Thus, our study provides a smaller but potentially important piece of evidence that humans may be able to learn how to distinguish between AI and non-AI; by developing a refined intuition for the often subtle differences between these opposites. Perhaps, in the future, it will become virtually impossible to develop a fully conscious understanding of such boundaries, but we may still be able to rely on our gut feelings. Our study suggests that this gut feeling, in a sense, is similar to other cognitive abilities in that it can be trained; and it may prove to be a useful skill across many different contexts. Therefore, we encourage future researchers to test the efficacy of the inductive learning paradigm in several different contexts and settings. For instance, future studies could assess the efficacy of inductive learning when it comes to training the recognition skills for AI generated text, photography, videos or other types of artworks. Research about the perception of AI is at a very early stage; therefore, we encourage scientists to engage in more exploratory research, which could stimulate the development of new theories in this domain.

When it comes to the artistic domain, there is plenty of room for speculation, even in opposing directions. On one hand, participants did not exhibit a better accuracy at guessing an artwork's origin than random chance; which may suggest that the perceptual boundary between AI and non-AI art has begun to fade away. On the other hand, the results suggested that participants, after training, could register patterns of "AI-ness", but not "human-ness" in the artworks – either unconsciously or consciously. Thus, somewhat paradoxically, perhaps the distinguishing feature of non-AI artworks is that they do not have features that they share in common; that the individuality of the creator is always present in the artwork. On the other hand, it was mentioned before that all the artworks in the study were generated by Midjourney; and it is possible that Midjourney has its own characteristic style. Thus, perhaps the patterns of 'AI-ness' are merely due to some specific features in the algorithms that they are based on; and perhaps each algorithm has its own 'artistic style'. In the light of these possibilities, future studies could compare artworks generated by different algorithms.

Overall, it seems that the artistic output of AI can indeed already highly resemble that of a human; and in the coming years, as AI technology will develop further, the outputs may become even more refined and differentiated. This, eventually, can call into question the very meaning of being an artist, or art itself, as well. For the most part of civilized history, art has been an integral part of humanity across all cultures; and artists have been highly respected individuals, often even deemed as visionaries. The existential or spiritual value of art may well be challenged by computers that reduce the creation of artistic objects to a set of mathematical formulas. Likewise, art has also carried considerable monetary value during much of history. As making art has never been faster and easier than with AI, this monetary value might shrink further with the inflation of the market – something that had already been happening in the music industry even before AI-music generation appeared (Brennan & Devine, 2020). This could potentially entail that many artists will not be able to make a living with their profession anymore; or that being an artist, both as a profession and as a socially valued status may weaken or disappear. In opposition though, these scenarios that predict AI dominance in art are not the only possibilities. Perhaps AI is able to create artworks that seem authentic, but that does not mean that it is already able to create art that is truly awe-inspiring. As our study did not focus on such evaluations, future studies could also investigate whether AI-art is able to inspire awe similar to some of the most renowned works of art; such studies could use the assistance of experts in art to choose the most suitable artworks for these purposes.

In addition, as AI will become more and more dominant in the art industry, it may also encourage artists to integrate AI art to their own artworks in unique ways, which may set human art further apart from AI art again. The usage of externally generated ideas has not been uncommon in the world of art even before the appearance of AI. In music production, for instance, generative synthesis has already been possible with analogue synthesizers for decades; and the role of the producer was, partly, to filter out elements that stand out from the rest and place those into the right context. This way of working may be applied in visual art as well: generated artworks might provide useful ideas which can be rearranged or combined with other original elements. Thus, it may also be the case that the existential and monetary concerns of artists will encourage them to reinvent themselves, which might also lead to fascinating transformations to the way humans express their artistic urges.

References

Abdukarimova, N. A., & Zubaydova, N. N. (2021). Deductive and inductive approaches to teaching grammar. *JournalNX*, 372-376.

Badet, J. (2021). AI, automation and new jobs. *Open Journal of Business and Management*, *9*(5), 2452-2463.

Bellaiche, L., Shahi, R., Turpin, M. H., Ragnhildstveit, A., Sprockett, S., Barr, N., ... & Seli, P. (2023). Humans versus AI: whether and why we prefer human-created compared to AI-created artwork. *Cognitive Research: Principles and Implications*, *8*(1), 42.

Betsch, T., & Glöckner, A. (2010). Intuition in judgment and decision making: Extensive thinking without effort. *Psychological Inquiry*, *21*(4), 279-294.

Birnbaum, M.S., Kornell, N., Bjork, E.L. *et al.* Why interleaving enhances inductive learning: The roles of discrimination and retrieval. *Mem Cogn* **41**, 392–402 (2013). https://doi.org/10.3758/s13421-012-0272-7

Brennan, M., & Devine, K. (2020). The cost of music. *Popular Music*, *39*(1), 43–65. doi:10.1017/S0261143019000552

Britannica. (n.d.). Why does AI art screw up hands and fingers? In Encyclopædia Britannica. Retrieved from https://www.britannica.com/topic/Why-does-AI-art-screw-up-hands-and-fingers-2230501

Dopkins, S., & Gleason, T. (1997). Comparing exemplar and prototype models of categorization. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, *51*(3), 212.

Epstein, Z., Levine, S., Rand, D. G., & Rahwan, I. (2020). Who gets credit for AI-generated art?. *Iscience*, *23*(9)

Ha, A. Y. J., Passananti, J., Bhaskar, R., Shan, S., Southen, R., Zheng, H., & Zhao, B. Y. (2024). Organic or Diffused: Can We Distinguish Human Art from AI-generated Images?. *arXiv preprint arXiv:2402.03214*.

Kang, S. H., & Pashler, H. (2012). Learning painting styles: Spacing is advantageous when it promotes discriminative contrast. *Applied Cognitive Psychology*, *26*(1), 97-103.

Kornell, N., & Bjork, R. A. (2008). Learning concepts and categories: Is spacing the "enemy of induction"?. *Psychological science*, *19*(6), 585-592.

Midjourney. (2023). *Midjourney* (V5) [Text-to-image model]. https://www.midjourney.com/

Miller, E. J., Steward, B. A., Witkower, Z., Sutherland, C. A. M., Krumhuber, E. G., & Dawel, A. (2023). AI Hyperrealism: Why AI Faces Are Perceived as More Real Than Human Ones. Psychological Science, 34(12), 1390-1403. https://doi.org/10.1177/09567976231207095)

Samo, A., & Highhouse, S. (2023). Artificial intelligence and art: Identifying the aesthetic judgment factors that distinguish human-and machine-generated artwork. *Psychology of Aesthetics, Creativity, and the Arts*.

Sedlmeier, P. (2014). From associations to intuitive judgment and decision making: Implicitly learning from experience. In *The routines of decision making* (pp. 83-99). Psychology Press.

uDiscoverMusic. (n.d.). Beethoven's 10th Symphony Reimagined by AI. Retrieved from https://www.udiscovermusic.com/classical-news/beethovens-10th-symphony-ai/

Zhou, Y., & Kawabata, H. (2023). Eyes can tell: Assessment of implicit attitudes toward AI art. *i-Perception*, *14*(5), 20416695231209846.

Appendix A

Tables and Figures

Figure 1

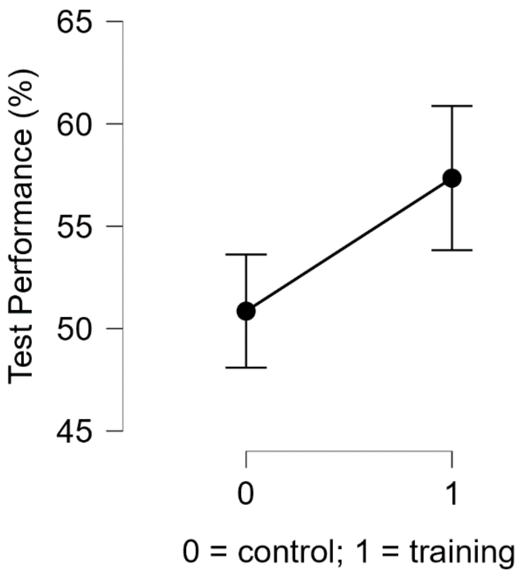

Table 1

Performances in the Control Group and Experimental Group

(Percentage Correct)

	Without Training	**With Training**
Valid	47	35
Missing	3	10
Mean	50.854%	57.348%
Std. Deviation	9.409	10.249
Minimum	30.952	31.250
Maximum	75.000	78.571

Table 2

Regression Coefficients for non-AI Art and AI Art Predicting Test Performance

Model	**Predictor**	**B**	**SE**	**β**	**t**	**p**
H_1	Intercept	53.575	1.761		46.631	< .001
	Art Affiliation	-0.750	0.953	-0.090	-0.787	0.434
	AI Art Affiliation	-0.274	1.007	-0.031	-0.272	0.787

Table 3

Regression Coefficients for Confidence Predicting Test Performance.

Model	**Predictor**	**B**	**SE**	**β**	**t**	**p**
H_1	Intercept	54.997	4.383		12.548	< .001
	Confidence	-0.023	0.072	-0.036	-0.324	0.747

Table 4

Mean Performances Across Types of Artworks (Percentage Correct)

	N	**Mean**	**SD**
Portrait	82	0.570	0.156
Landscape	82	0.539	0.135
Abstract	82	0.501	0.132

Appendix B

AI picture creation prompts

Abstract

abstract, oil on canvas painting like Max de Winter's Monkey Business, that is seemingly unstructured at first glance, but does have from human-ish shapes that come together in the strokes, which are not too obvious. the colours should be a little bit darker than in the original work.

multiple abstract modern paintings

multiple abstract modern paintings

multiple abstract modern paintings

abstract oil on canvas paining in the style of abreesha jones, using the same brushes as the artist does.

abstract oil on canvas painting in the style of lisa carney. use the same paining technques and brushes as the artist

oil on canvas paninting exactly like this but with slightly different shapes and arrangement

abstract but realistically structured, oil on canvas painting that seems to resemble a futuristic, dystopian, but slightly humorous city. sophisticated use of brush and strokes

abstract painting of intertwined zebra's filling up the entire frame only in black and white, figurative, victor vasarely

Agamograph by Yaacov Agam

an abstract painting

an abstract painting

an abstract painting

an abstract painting

an abstract painting

an abstract painting

an abstract painting

Homage to the Square by Josef Albers

minimalistic abstract painting in this style, without any shapes of humans or anything figurative. should suggest the feeling of falling apart

simple, abstract painting, using different shades of orange, also playing with the strenght of pushing the brush against the canvas. and simple repeating patterns of hexagons, in a neat, simple arrangement. should represent the feeling of coming together.

Landscape

Simon Stålenhag

oil on canvas landscape in sunrise, depicting a flat, clean but picturesque field. in the style of Richmond Castle.

a landscape painting that looks like meindert hobbema's work

a landscape painting that looks like meindert hobbema's work

a landscape painting that looks like meindert hobbema's work

a landscape painting that looks like peter paul rubens' work.

april gornik dunes behind savanna monotonous sky

april gornik dunes behind savanna monotonous sky

april gornik dunes behind savanna monotonous sky

a landscape painting that looks like peter paul rubens' work.

a landscape painting that looks like peter paul rubens' work.

erin hanson cherry blossom

erin hanson arbor of light

oil on canvas landscape after sundown, with a vibrant, purple, but still realistic sky, depicting a slightly hilly, clean but picturesque field. in the style of Herman van Swanevelt.

oil on canvas landscape in sunrise, depicting a flat, clean but picturesque field. in the style of Richmond Castle.

april gornik wheatfield with monotonous dark sky and a tree

fine brush painting on canvas in the style of Jacob van Ruisdael, depicting a river, rocks, and a small waterfall.

fine brush painting on canvas in the style of Jacob van Ruisdael, depicting a river, rocks, and a small waterfall.

a brush painting of a rainy dutch forest and farmland, in the style of peter paul rubbens.

a brush painting of a rainy dutch forest and farmland, in the style of peter paul rubbens.

Portraits

portrait 18th century rococo neoclassicism grand manner chiaroscuro sfumato pastoral patronage allegory physiognomy gaze drapery vanitas face

francisco de goya

create an oil portrait of John the baptist using the alla prima painting technique on canvas make sure that the face is painted using the underpainting technique

create an oil portrait of marie antoinette using the alla grisaille technique on canvas make sure that the face is painted using the underpainting technique

create an oil portrait of John the baptist using the impasto technique on canvas make sure that the face is painted using the underpainting technique

create a full body portrait of John the Baptist in front of the Jordan River using the alla prima technique on canvas, make sure that the face is painted using the underpainting technique

create a full-body oil portrait of Moses holding the Ten Commandments using the alla prima technique on canvas, and make sure that the face is painted using the underpainting technique

create a full-body oil portrait of Moses holding the Ten Commandments written on stone tablets in an impressionist style using the alla prima technique on canvas, and make sure that the face is painted using the underpainting technique

a baroque style oil on paint portrait of a merchant

paint a portrait of a merchant, standing in front of cart, using oil paints on canvas and the impasto painting technique

a baroque-style oil on canvas portrait of a monk

old renaissance portrait of a 14th century peasant

old renaissance portrait of a 15th century wealthy man

a portrait painting, that looks like Rembrandt's work painted portrait old dark canvas oil beggar

a dark and old-fashioned full-body oilpainting on canvas portrait of Anne Hutchinson in the style of Rembrandt

a dark and old-fashioned full-body oilpainting on canvas portrait of Anne Hutchinson in the style of Rembrandt

old renaissance portrait of a wealthy merchant 15th century

paint a baroque-style oil on canvas portrait of a young, 17th-century princess standing in an orchard

paint a baroque-style oil on canvas portrait of a young, 17th-century princess standing in an orchard

Human-made picture titles

Abstract

Orange Blossom-Lisa Carney

Homage to the square- Joseph Albers

Healing Antenna- Matthew Dibble

Monkey business- Max de Winter

Told you so!- Max de Winter

The Trendsetter- Max de Winter

Typografie Design- Henry Stazewski

Relief- Henry Stazewski

Vicky Barranguet- All about you

Jeffrey Tover- Coachella Valley

Naomi Yuki- Cosmos, Inside

Victor Vasarely- Zebras

Sonia Delaunay- Electric Prisms

Jeffery Tover- Los Angeles

Jeffrey Tover- Night Ride

Vicky Barranguet- Nothing held back

Vicky Barranguet- Roads not taken

Paul Franklin- Turquoise Moon

Kazimir Malevich- Dynamic Suprematism

Landscape

Haystacks: Autumn - Jean-Francois Millet

Landscape Study with Clouds - Emile Loubon

Cuckmere Haven - Eric Ravilious

Grainfields - Jacob van Ruisdeel

Landscape by Moonlight - Peter Paul Rubens

Landscape - Circle of Carl Rottmann

Mountainous Landscape at Vicovaro - Simon Denis

The Waterspout - Gustave Courbet

View of Tivoli from Santa Maria del Giglio - Leon Fleury

The Alley at Middelharnis - Meindert Hobbema

Meindert Hobbema- Watermolen

Achille Etna Michallon- Waterfall at Mont-Dore

Eugene Isabey- Sunset on the Normandy Coast

Simon Denis- On the Quirinal Hill

R.S. Duncan- Savanna

Philip Wilson Steer- Richmond Castle, Yorkshire

Eric Hanson- Cherry Blossom

Simon Stalenhag- The Mascot

Claude Lorrain- Sunrise

Paul Cezanne- Viaduct of the Arc River Valley

Portrait

Portrait of an Unknown Woman - Ivan Kramskoy

Jean-Baptiste Faure - Edouard Manet

Reading Woman - Ivan Kramskoy

Comtesse de la Châtre - Élisabeth Vigée Le Brun

Archbishop of Milan - Tiziano Vecellio

Portrait of Dmitri Vasilievich Grigorovich - Ivan Kramskoy

Francois Gerard - Antoine-Jean Gros

FLINT OIL ON LINEN 2017 (MISSING)

The Love Letter - Jean-Honore Fragonard

Samuel P. Avery - Raimundo de Madrazo y Garreta

Portrait of a man - Unknown artist

Lady Elizabeth Stanley - George Romney

Portrait of Louis-Félix Amiel - Eugène Devéria

Lucia - Frederic Leighton

Portrait of a Man - David Bailly

Portrait of Claes Duyst van Voorhout - Frans Hals

Sibylle - Corot

Marie Joséphine Charlotte du Val d'Ognes - Marie Denise Villers

Mrs. Richard Bache - John Hoppner

Portrait of a Child - Camille Corot

Vienna Art Interest and Knowledge Questionnaire (VAIAK) (Specker et al., 2020)

3. I like to talk about art with others.

7. I'm interested in art.

9. I'm always looking for new artistic impressions and experiences.

10. In everyday life I routinely see art objects that fascinate me.

12. How often do you visit art museums and/or galleries?

13. How often do you read books, magazines or catalogs about art?

14. How often do you look at images of artworks (catalogs, internet, etc.)?

Vienna Art Interest and Knowledge Questionnaire (VAIAK) (Specker et al., 2020), Adapted AI Scale

1. I like to talk about AI art technology with others.

2. I’m interested in AI art technology.

3. I’m always looking for new AI art impressions and experiences.

4. In everyday life I see AI art that fascinates me.

5. How often do you seek out AI art technology?

6. How often do you read articles about AI art technology?

7. How often do you look at AI artwork and images (e.g. on the internet, etc.)?